PBS Job Queue Structure

To run a PBS job on the Pleiades, Aitken, or Electra compute nodes, submit the job from a Pleiades front-end system (PFE) to one of the queues managed by the PBS server pbspl1. To view a complete list of all queues, including their defaults or limits on the number of cores, wall time, and priority, use the qstat command as follows:

pfe*21*% qstat -Q

To view more detailed information about each queue, such as access control, defaults, limits on various resources, and status:

pfe*21*% qstat -fQ *queue_name*

Brief summaries of the various queues are provided in the following sections.

Queues for Pleiades, Aitken, and Electra Users

You can run jobs on Pleiades, Aitken and Electra using any of the following queues:

```
Queue   NCPUs/      Time/
name     max/def    max/def    pr
======= =====/=== ======/====== ===
low        --/  8  04:00/ 00:30 -10
normal     --/  8  08:00/ 01:00   0
long       --/  8 120:00/ 01:00   0
debug      --/  8  02:00/ 00:30  15
devel      --/  1  02:00/    -- 149
```

The normal, long and low queues are for production work. The debug and devel queues have higher priority and are for debugging and development work.

For the debug queue, each user is allowed a maximum of 2 running jobs using a total of up to 128 nodes. For thedevel queue, each user is allowed only 1 job (either queued or running).

See the following articles for more information:

- Pleiades devel Queue
- Preparing to Run on Electra Broadwell Nodes
- Preparing to Run on Electra Skylake Nodes
- Preparing to Run on Aitken Cascade Lake Nodes

Queues for GPU Users

Jobs submitted to the v100 queue are routed to run on a separate PBS server, pbspl4, for nodes with four or eight NVIDIA Volta V100 GPU cards. See the following articles for more information:

- Requesting GPU Resources
- Changes to PBS Job Requests for V100 GPU Resources

Queues for Endeavour Users

Jobs submitted from PFEs to the e_normal, e_long, e_debug, and e_low queues are routed to run on the Endeavour shared-memory system, which has its own PBS server. Access to these queues requires allocation on Endeavour.

See Endeavour PBS Server and Queue Structure for more information.

Queue for Lou Users

The ldan queue provides access to one of the Lou data analysis nodes (LDANs) for post-processing archive data on Lou.

The ldan queue is managed by the same PBS server as all of the Pleiades queues and is accessible by submitting jobs from an LFE or PFE.

Each job on the ldan queue is limited to 1 node and 72 hours. Each user can have a maximum of 2 running jobs on this queue.

See Lou Data Analysis Nodes for more information.

Queues with Restricted Access

The vlong Queue

Access to the vlong queue requires authorization from the NAS User Services manager.

A maximum of 1024 nodes and 16 days are allowed on this queue.

To avoid losing all useful work in the event of a system or application failure, ensure that your application includes checkpoint and restart capability; this is especially important when using this queue.

The wide Queue

The wide queue is a high priority queue used for jobs that meet either one of the following conditions:

- The number of requested NCPUs (number of nodes * number of cores per node) exceeds 25% of the user's mission shares
- The number of nodes requested of a processor model exceeds 25% of the total number of nodes of that model. See the Pleiades, Electra, and Aitken configuration pages for the total number of nodes for each model.

Users do not have direct access to the wide queue. Instead, a job that meets the above criteria is moved into the wide queue by NAS staff. Only 1 job in the wide queue is allowed among all HECC systems (Pleiades, Electra, and Aitken). When more than one job meets the criteria, preference is given to the job that has been waiting for longer time or whose owner does not have other jobs running.

Various Special Queues

Special queues are higher priority queues for time critical projects to meet important deadlines. Some examples of special queues are armd_spl, astro_ops, heomd_spl, kepler, smd1, smd2, smd_ops, and so on. Request for creation and access of a special queue is carefully reviewed by the mission directorate or by the program POC who authorizes the project's SBU allocation. The request is rarely granted unless the project's principal investigator demonstrates that access to a special queue is absolutely necessary.

Advanced Reservation Queues

An advanced reservation queue allows you to reserve resources for a certain period ahead of time. One circumstance in which the NAS User Services manager may authorize the use of a reservation queue is when a user plans to run an extra-large job, such as one exceeding 15,000 cores, which would never start running without a reservation.

A reservation queue is created by NAS staff for the user and is normally named with the letter R followed by a number.

To find the status of existing PBS advanced reservations, issue the pbs_rstat command on a PFE.

Note: Queues such as alphatst, ded_time, diags, dpr, idle, resize_test, and testing are used for staff testing purposes, and are not accessible by general users.

Article ID: 187
Last updated: 31 Aug, 2022
Updated by: Moyer M.
Revision: 56
Running Jobs with PBS -> PBS on Pleiades -> PBS Job Queue Structure
https://www.nas.nasa.gov/hecc/support/kb/entry/187/